

REMARKS

The Office Action dated February 25, 2009, has been received and carefully noted. The above amendments to the claims, and the following remarks, are submitted as a full and complete response thereto.

Status of the Claims

Claims 1, 9, 18, 24, 30 and 31 have been amended to more particularly point out and distinctly claim the subject matter of the invention. No new matter has been added. Claims 1-9, 11-20, 22-25 and 27-34 are currently pending in the application and are respectfully submitted for consideration.

Rejection under 35 U.S.C. § 103

Claims 1-9, 11-20, 22-25 and 27-34 were rejected under 35 USC § 103(a) as allegedly being unpatentable over Bruck et al. (U.S. Patent No. 6,691,165) in view of Sato (U.S. Patent No. 7,277,935) and further in view of McLaughlin et al. (U.S. Publication No. 2002/0165929). The Office Action took the position on pages 2-7 that the combination of Bruck et al., Sato and McLaughlin et al. teaches all of the features of the rejected claims. Applicants respectfully traverse the rejection. Reconsideration of the claims is respectfully requested.

Independent claim 1, from which claims 2-8 and 27 depend, recites a system including a network interface configured to communicate with nodes in a cluster and a configuration subsystem operationally coupled to a remote management broker. The remote management broker is configured to distribute information between the nodes in

the cluster. The system also includes a processor configured to access the cluster from a single-point, obtain information relating to at least two devices within the cluster, present the information to a user and determine network management operations to perform on the cluster. The processor is also configured to apply a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and to perform the determined network management operations and to determine whether the network management operations on the cluster, including the at least two devices, were applied correctly. When the network management operations were not applied correctly, the processor is configured to roll back to the most recent successful configuration.

Independent claim 9, from which claims 11-17 and 28 depend, recites a method including accessing a cluster from a single-point, obtaining attributes relating to at least two devices within the cluster, presenting the attributes to a user and receiving input from the user relating to the attributes and determining network management operations to perform on the cluster based on the received input. The method also includes applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and performing the determined network management operations on the cluster. The method further includes determining whether the network management operations on the cluster, including the at least two devices, were applied correctly. When the network management operations were not

applied correctly, the method further includes rolling back to the most recent successful configuration.

Independent claim 18, from which claims 19, 20, 22, 23 and 29 depend, recites a computer program embodied on a computer readable storage medium configured to control a processor to perform a process, including accessing a cluster from a single-point, obtaining attributes relating to at least two devices within a cluster from a single-point and presenting the attributes to a user and receiving input relating to the attributes from the user. The process also includes determining network management operations to perform on the cluster based on the received input, applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and performing the determined network management operations on the cluster. The process further includes determining whether the network management operations on the cluster, including the at least two devices, were applied correctly, and when the network management operations were not applied correctly, rolling back to the most recent successful configuration.

Independent claim 24, from which claim 25 depends, recites an apparatus including obtaining means for obtaining attributes relating to at least two devices within a cluster from a single-point, providing means for providing the attributes to a user, receiving means for receiving input relating to the attributes from the user, determining means for determining network management operations to perform on the cluster based on the received input and distributing means for distributing the network management

operations to the devices within the cluster. The apparatus also includes first applying means for applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and second applying means for applying the network management operations to the devices within the cluster. The apparatus further includes determining means for determining whether the network management operations on the cluster, including the at least two devices, were applied correctly, and when the network management operations on the cluster were not applied correctly, rolling back to the most recent successful configuration.

Independent claim 30 recites an apparatus including accessing means for accessing the cluster from a single-point, obtaining means for obtaining attributes relating to at least two devices within the cluster and presenting means for presenting the attributes to a user and receiving input from the user relating to the attributes. The apparatus also includes operation determining means for determining network management operations to perform on the cluster based on the received input, applying means for applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and performing means for performing the determined network management operations on the cluster. The apparatus further includes correction determining means for determining whether the network management operations on the cluster, including the at least two devices, were applied correctly, and when the network management operations were not applied correctly, rolling back to the most recent successful configuration.

Independent claim 31, from which claims 32-34 depend, recites an apparatus including a processor configured to access a cluster from a single-point, obtain attributes relating to at least two devices within the cluster and present the attributes to a user and receive input from the user relating to the attributes. The processor is also configured to determine network management operations to perform to the cluster, apply a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and perform the determined network management operations. The processor is further configured to determine whether the network management operations on the cluster, including the at least two devices, were applied correctly, and when the network management operations were not applied correctly, the processor is configured to roll back to the most recent successful configuration.

As will be discussed below, Bruck et al., Sato and McLaughlin et al., both individually and in combination, fail to teach or suggest all of the features of the presently pending claims.

Bruck et al. generally discusses “a scalable, distributed, highly available, load balancing server system having multiple machines functioning as a front server layer between the network and a back-end server layer having multiple machines functioning as Web file servers, FTP servers, or other application servers” (column 2, lines 42-47). “The operation of the servers on both layers is monitored, and when a server failure at either layer is detected, the system automatically shifts network traffic from the failed

machine to one or more operational machines, reconfiguring front-layer servers as needed without interrupting operation of the server system” (column 2, lines 49-54, of Bruck et al.). “The front layer machines perform their operations without breaking network communications between clients and servers, and without rebooting of computers. In this way, the system [allegedly] provides reliable network communication in a scalable load balancing solution for server farms” (column 2, lines 63-67, of Bruck et al.).

Sato generally discusses “management of a communications apparatus” and the discussion thereof is allegedly applicable “to initial settings by a management device on a network device such as a hub and a router connected to a network” (see column 1, lines 7-11). A method as discussed in Sato may include:

providing, via one or more interconnecting devices connected to a network, an address of a target device which is connected to one of the interconnecting devices, or a port identifier for identifying a port of the target device which is connected to one of the interconnecting devices via a port of one of the interconnecting devices, for a management device that is connected to the network so as to manage the target device; and the management device managing the target device based upon the provided address. According to this management method, the management device manages the target device based upon the address or port identifier provided from the interconnecting device. This allows the address or port identifier to be automatically provided from the interconnecting device to the management device, and thus facilitates the management including initial configuration.

(Column 3, lines 1-17).

McLaughlin et al. generally discusses “the allocation, retention, and release of control of a programmable switch in a switched network to allow maintenance to be performed thereon” (paragraph [0005]). “[A] mechanism is provided that allows for

synchronization of online maintenance of the firmware and configuration files in the programmable switch components of a network. At the same time, the ability to direct that maintenance remains distributed among the nodes of the network” (paragraph [0012], of McLaughlin et al.).

Independent claim 1 recites, in part, that “when the network management operations were not applied correctly, the processor is configured to roll back to the most recent successful configuration.” Independent claims 9, 18, 24, 30 and 31, which each have their own scope, recite similar features. The Office Action conceded on page 3 that Bruck et al. does not disclose “determining whether the network management operations on the cluster, including said at least two devices, were applied correctly, and when the network management operations were not applied correctly, the processor is configured to roll back to a successful configuration.” Rather, the Office Action relied on column 16, lines 10-17, of Sato to allegedly cure these deficiencies. Applicants respectfully submit that the cited art fails to teach or suggest these features.

The cited section of Sato discusses that:

TFTP transmits and receives mainly configuration information and firmware of the network device 60 in the form of a file. For example, the configuration information is collected as a file, and if the network device 60 fails, the file is transmitted by return and used when the original configuration information should be restored. Under HTTP, HTML (Hypertext Markup Language) documents are transmitted and received between the management device 10 and the network device 60.

Thus, in Sato, a non-clustered management device may transmit a file to a network device that is used to restore the **original** configuration information in the firmware of,

for example, a network switch. On the other hand, claim 1 recites that the processor is configured to roll back to the **most recent successful configuration**. In a dynamic cluster computing environment, the original configuration will likely be quickly outdated, and perhaps outdated with an updated configuration used for the initial deployment of the cluster. Only restoring to an original configuration, as discussed in Sato, clearly fails to teach or suggest these features.

As noted in the previous Response filed February 4, 2009, as discussed on page 1, lines 24-27, of the present specification, “[o]ne problem is that it is difficult to maintain identical configurations of the Network Management features on all devices within a cluster. In addition, errors in the configuration of one device, or incompatible configurations among the devices, may render a particular NM feature inoperable.” By rolling back to a successful configuration, some embodiments of the present invention have the advantage of being able to maintain successful configurations that occurred **after** the initial implementation of the cluster, which allows better management and adaptability of the cluster throughout its lifetime. Such features are simply absent from, and such advantages simply cannot be achieved with, the system discussed in Sato. Further, nothing is cited or found in McLaughlin et al. that overcomes the above deficiencies of Bruck et al. and Sato.

Applicants also respectfully submit that the combination of Sato with Bruck et al. and McLaughlin et al. is improper. The Office Action argued on page 8 that “as evidence[d] by the fact that all three prior art references are classified in the same field of

endeavor (i.e. 709 ‘multicomputer data transferring’), it would suggest that the teachings of Sato would not have been outside the scope of a person of ordinary skill in the art of computer networking even though, as suggested by the applicant, Sato may not use the term ‘cluster’ explicitly in his reference.” However, Applicants submit that “multicomputer data transferring” has an extremely broad scope as nearly every networked computing device transfers some form of data to other networked computing devices. The scope is so broad that it would not be useful to one of ordinary skill in the art to determine which art would potentially be a successful candidate to provide a potential solution to certain problems in the field of cluster computing. In fact, the Examiner appeared to recognize the difficulty of searching the entirety of class 709 and instead narrowed the search to subclasses 212-216 and 223-229 (see enclosed Examiner’s Search Notes). The field of cluster computing, on the other hand, is much more appropriate for the present application, which is entitled “Single-Point Management System for Devices in a Cluster”.

Applicants also submit that the application of *KSR* in the Office Action is improper. As the Office Action noted, *KSR International Co. v. Teleflex Inc.* stated that “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement predictable variations, §103 likely bars its patentability” (550 U.S. 398, 411 (2007)). The Office Action alleged on page 8 that “Sato teaches an implicit design incentive in restoring a network device to an original configuration (see

col. 16, line[s] 10-15), in that it allows a failed device to continue operations using an original configuration in a ‘non-clustered system’. Furthermore, there does not appear to be any technological barrier that would have prevented a person of ordinary skill in the art from applying the teachings of Sato to a ‘clustered system’ as taught by McLaughlin and Bruck.”

However, such a technical barrier clearly exists, and it further has not been demonstrated that the original configuration restoration operations of Sato would work in a system where the configurations of clustered devices are managed. Per the above, one problem is that it is difficult to maintain **identical** configurations of the network management (NM) features on **all devices** within a cluster. In addition, errors in the configuration of one device, or **incompatible configurations among the devices**, may render a **particular NM feature** inoperable. Merely adding a means of restoring a device to an original configuration, as discussed in the non-clustered Sato system, would be incompatible with a practical cluster management implementation. Rolling back to an initial configuration for a device could, for instance, render a subsequently implemented NM feature inoperable. Further, the cited art offers no guidance as to how rolling back configurations could be coordinated among the multiple devices of a cluster. Thus, the combination of Bruck et al., Sato and McLaughlin et al. is not a “predictable use of prior art elements according to their established functions”, as required by *KSR*. The claimed features are interdependent and must be read as a whole, and the addition of Sato in the manner proposed in the Office Action would destroy practical implementations of some

embodiments of the system recited in claim 1. Thus, *KSR* not only fails to support, but explicitly rejects, the combination of Sato with Bruck et al. and McLaughlin et al.

Claims 2-8, 11-17, 19, 20, 22, 23, 25, 27-29 and 32-34 depend from independent claims 1, 9, 18, 24 or 31 and add further features thereto. Thus, the arguments above with respect to the independent claims also apply to the dependent claims.

Per the above, Bruck et al., Sato and McLaughlin et al., both individually and in combination, fail to teach or suggest all of the features of the above-rejected claims under 35 U.S.C. § 103(a). Accordingly, it is respectfully submitted that the rejection is overcome and respectfully requested that the rejection be withdrawn.


Conclusion

For at least the reasons presented above, it is respectfully submitted that claims 1-9, 11-20, 22-25 and 27-34, comprising all of the currently pending claims, patentably distinguish over the cited art. Accordingly, it is respectfully requested that the claims be allowed and the application be passed to issue.

If for any reason the Examiner determines that the application is not now in condition for allowance, it is respectfully requested that the Examiner contact, by telephone, Applicants' undersigned representative at the indicated telephone number to arrange for an interview to expedite the disposition of this application.

In the event this paper is not being timely filed, Applicants respectfully petition for an appropriate extension of time. Any fees for such an extension together with any additional fees may be charged to Counsel's Deposit Account 50-2222.

Respectfully submitted,




Michael A. Leonard II
Attorney for Applicants
Registration No. 60,180

Customer No. 32294
SQUIRE, SANDERS & DEMPSEY LLP
14TH Floor
8000 Towers Crescent Drive
Vienna, Virginia 22182-6212
Telephone: 703-720-7800
Fax: 703-720-7802

MAL:jf

Enclosure: Examiner's Search Notes

Search Notes 	Application/Control No. 10603925	Applicant(s)/Patent Under Reexamination TREPPA ET AL.
	Examiner Higa, Brendan	Art Unit 2153

SEARCHED			
Class	Subclass	Date	Examiner
709	212-216, 223-229	12/22/2008	BH

SEARCH NOTES		
Search Notes	Date	Examiner
see EAST search history	12/22/2008	BH

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner

--	--